

ABSTRACT OF THE DISCLOSURE

To automatically evaluate an essay for overly repetitive word usage, a word is identified in the essay and at least one feature associated with the word is determined. In addition, a probability of the word being used in an overly repetitive manner is determined by mapping the feature to a model. The model having been generated by a machine learning application based on at least one evaluated essay. Furthermore, the essay is annotated to indicate the word is used in an overly repetitive manner in response to the probability exceeding a threshold probability.